arXiv:1402.5431v1 [math.ST] 21 Feb 2014

# Comparing two formulations of skew distributions with special reference to model-based clustering

Adelchi Azzalini* Ryan P. Browne†
Marc G. Genton‡ Paul D. McNicholas†

**Abstract**

Multivariate skew distributions continue to gain popularity as effective tools in data analysis for a wide variety of data types. They also feature in several modern computational statistics techniques, including model-based clustering approaches. Two well-known formulations of skew distributions have been used extensively within the model-based clustering literature over the past few years. We investigate the properties of these formulations and, in doing so, we refute various claims that have been made about these formulations of late. Our position as to the inaccuracy of these claims is supported by theoretical arguments as well as real data examples.

## 1 Background and Motivation

In the last 10–15 years, much work in the statistical literature has been dedicated to the development of flexible parametric classes of distributions, with special interest in the multivariate case where the set of alternatives to the normal distribution is not as wide as in the univariate case. When the domain of interest is the $d$-dimensional Euclidean space, the class of elliptical distributions is the more studied alternative to the normal family, of which it represents a superset. However, because the elliptical class imposes stringent symmetry conditions, which must hold simultaneously for all components of the distribution, this calls for the development of more flexible families to more closely meet the data behaviour in a number of real cases.

When one moves away from the symmetry of the multivariate normal or the elliptical distributions, the feature that arises more easily is skewness, and this explains the widespread use of the prefix 'skew' which recurs almost constantly in this context. The specific, but already quite broad, stream of literature we are referring to is commonly identified by terms like 'skew-symmetric', 'skew-elliptical' and similar ones. A recent extensive account is provided by Azzalini and Capitanio (2014).

This activity has generated an enormous number of formulations, sometimes arising with the same motivation and target, or nearly so. A natural question in these cases is which of the competing alternatives is preferable, either universally or for some given purpose. To be more specific, start by considering the multivariate skew-normal (SN) distribution proposed by Azzalini and Dalla Valle (1996), examined further by Azzalini and Capitanio (1999) and by much subsequent work. Note that, although the latter paper adopts a different parameterization of the earlier one, the set of distributions that they encompass is the same; we shall denote this construction as the 'classical skew-normal'. Another form of skew-normal distribution has been studied by Sahu, Dey, and Branco (2003), which we shall refer to as the 'SDB skew-normal'. The classical and the SDB set of distributions coincide only for dimension $d = 1$; otherwise, the two sets differ and not simply because of different parameterizations. For $d > 1$, the question then arises about whether

*Department of Statistical Science, University of Padua, 35121 Padova, Italy.
†Department of Mathematics & Statistics, University of Guelph, Guelph, Ontario, Canada.
‡CEMSE Division, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia.

there is some relevant difference between the two formulations from the viewpoint of suitability for statistical work, both on the side of formal properties and on the side of practical work.

The above-described situation carries on with the broader formulations that arise when in the underlying parent distribution the normal family is replaced by the wider elliptical class, leading to the so-called skew-elliptical distributions. A special case that has received much attention is the skew-$t$ family, which was formulated by Branco and Dey (2001) and examined in detail by Azzalini and Capitanio (2003). Again, this 'classical' skew-$t$ has a counterpart given by another skew-$t$ considered by Sahu et al. (2003), and the same questions as above hold.

For both forms of skew-$t$ distribution, a major reason for their appeal is represented by the great flexibility of the density shape, which arises because we can regulate the tail weight via the degrees of freedom, like with the traditional Student's $t$ distribution, as well as the slant of the distribution via a $d$-vector asymmetry parameter. A substantial fraction of the work in this direction has been generated in the context of model-based clustering, which involves introducing a mixture density of the form

$$f(x) = \sum_{g=1}^{G} \pi_g f_0(x; \theta_g), \qquad x \in \mathbb{R}^d, \tag{1}$$

where $\pi_g > 0$, with $\sum_g^G \pi_g = 1$, is the mixing proportion for the $g$th component, $f_0$ is a parametric family of distributions and $\theta_g$ is the parameter set of the $g$th mixture component within this family. The additional flexibility of the skew-$t$ distribution if used in the role of $f_0$, compared to the multivariate normal and Student's $t$, can diminish the number of components $G$ required to fit the data adequately when the data points are grouped in non-elliptical clusters, leading to a more easily parsimonious and a possibly more interpretable outcome.

There is an additional reason for paying special attention to this portion of the model-based clustering literature. In connection with the use of the above-recalled distributions, some views have been formulated quite recently which we believe to be unjustified and inaccurate, with effects that might potentially have unintended impact beyond the model-based clustering community. We are referring to the formulation of Lee and McLachlan (2014), to which for brevity we shall indicate as LM2012 because the paper appeared online in 2012. The authors have subsequently adopted the same framework in other publications (e.g., Lee and McLachlan, 2013b,c).

The overall aim of the present paper is to present a detailed comparison of the two formulations, classical and SDB, as for their formal properties and suitability in applied work, and to discuss the view formulated by Lee and McLachlan in the publications mentioned above.

# 2 Skew-normal and Skew-$t$ Distributions

## 2.1 Two skew-normal formulations

Because of their role as the basic constituent for more elaborate formulations, we start by discussing the two forms of skew-normal distributions. The density and the distribution function of a $N_d(0, \Sigma)$ variable are denoted $\varphi_d(\cdot; \Sigma)$ and $\Phi_d(\cdot; \Sigma)$, respectively; the $N(0, 1)$ distribution function is denoted $\Phi(\cdot)$.

The classical skew-normal density function is

$$f_c(x) = 2\, \varphi_d(x - \xi; \Omega)\, \Phi\{\alpha^\top \omega^{-1}(x - \xi)\}, \qquad x \in \mathbb{R}^d, \tag{2}$$

with parameter set $(\xi, \Omega, \alpha)$. Here $\xi$ is a $d$-dimensional location parameter, $\Omega$ is a symmetric positive definite $d \times d$ matrix, playing the role of scale parameter, $\alpha$ is a $d$-dimensional slant parameter and $\omega$ is a diagonal matrix formed by the square roots of the diagonal elements of $\Omega$.

Various stochastic representations exist for (2). One is as follows: if

$$\begin{pmatrix} X_0 \\ X_1 \end{pmatrix} \sim N_{d+1}(0, \Omega^*), \qquad \Omega^* = \begin{pmatrix} \bar{\Omega} & \delta \\ \delta^\top & 1 \end{pmatrix}$$

where $\Omega^*$ is a correlation matrix, then

$$Y_c = \xi + \omega(X_0 | X_1 > 0) \tag{3}$$

has distribution (2) with $\Omega = \omega\bar{\Omega}\omega$ and $\alpha = (1 - \delta^\top \bar{\Omega}^{-1}\delta)^{-1/2}\, \bar{\Omega}^{-1}\delta$. Another stochastic representation is the following: if $\delta$ is a $d$-vector with elements in $(-1, 1)$, (2) is the density function of

$$Y_c = \xi + \omega \left\{ [I_d - \mathrm{diag}(\delta)^2]^{1/2}\, V_0 + \delta |V_1| \right\}, \tag{4}$$

where $V_0$ and $V_1$ are independent normal variates of dimension $d$ and 1, respectively, with 0 mean value, unit variances and $\mathrm{cor}(V_0)$ is suitably related to $\alpha$ and $\Omega$.

The moment generating function of (2) is $M_c(t) = 2\exp(t^\top \xi + t^\top \Omega t/2)\Phi(t^\top \omega\delta)$, from which we obtain the mean value and the variance matrix

$$E(Y_c) = \xi + \sqrt{2/\pi}\, \omega\delta, \qquad \mathrm{var}(Y_c) = \Omega - (2/\pi)\, \omega\delta\, \delta^\top \omega,$$

where $\delta = (1 + \alpha^\top \bar{\Omega}\alpha)^{-1/2}\, \bar{\Omega}\alpha$. A more detailed exposition of the above facts is presented in Chapter 5 of Azzalini and Capitanio (2014), together with additional properties of the distribution.

For the SDB skew-normal, we adopt a very minor change from the symbols of Sahu et al. (2003), but retain the same parameterization. Given real values $\lambda_1, \ldots, \lambda_d$, let

$$\lambda = (\lambda_1, \ldots, \lambda_d)^\top, \qquad \Lambda = \mathrm{diag}(\lambda)$$

and write the SDB density as

$$f_s(x) = 2^d\, \varphi_d(x - \xi; \Delta + \Lambda^2)\, \Phi_d\{\Lambda(\Delta + \Lambda^2)^{-1}(x - \xi); I_d - \Lambda(\Delta + \Lambda^2)^{-1}\Lambda\}, \tag{5}$$

where $\Delta$ is a symmetric positive-definite matrix.

This density is associated with the following stochastic representation. For independent variables $\varepsilon \sim N_d(\xi, \Delta)$ and $Z \sim N_d(0, I_d)$, consider the transformation

$$Y_s = \Lambda(Z | Z > 0) + \varepsilon, \tag{6}$$

where $Z > 0$ means that the inequality is satisfied component-wise; then $Y_s$ has density (5). The moment generating function of $f_s$ is

$$M_s(t) = 2^d\, \exp\left\{ t^\top \xi + t^\top (\Delta + \Lambda^2) t/2 \right\} \prod_{j=1}^{d} \Phi(\lambda_j t_j), \tag{7}$$

which by differentiation leads to

$$E(Y_s) = \xi + \sqrt{2/\pi}\, \lambda\,, \qquad \mathrm{var}(Y_s) = \Delta + (1 - 2/\pi)\Lambda^2\,. \tag{8}$$

To obtain a first perception of the density shape of the two distributions, we produce plots which are the direct analogue of Figures 4 and 5 of Azzalini and Dalla Valle (1996) originally used to illustrate the classical skew-normal. The left panel of our Figure 1, which matches their Figure 4, displays a scatter plot of the weight (Wt, in kg) and the height (Ht, in cm) of 202 athletes examined at the Australian Institute of Sport (AIS), with the classical SN fit superimposed; in addition, we also plot the corresponding fit of the SDB

density. The right panel of Figure 1 displays a similar plot for the lean body mass (LBM) and body mass index (BMI) of the same athletes. The contour level curves of the classical and the SDB densities turn out to be very similar. In the left panel, the two sets of contour curves are virtually identical; in the right panel, the modes are somewhat displaced and the SDB curves have more spiky corners than the classical density, but again the densities are otherwise very similar. The AIS data, which comprise several other variables, have been used extensively for illustrative purposes in the subsequent literature, effectively becoming a sort of standard benchmark; therefore, we shall employ them again in the coming sections.

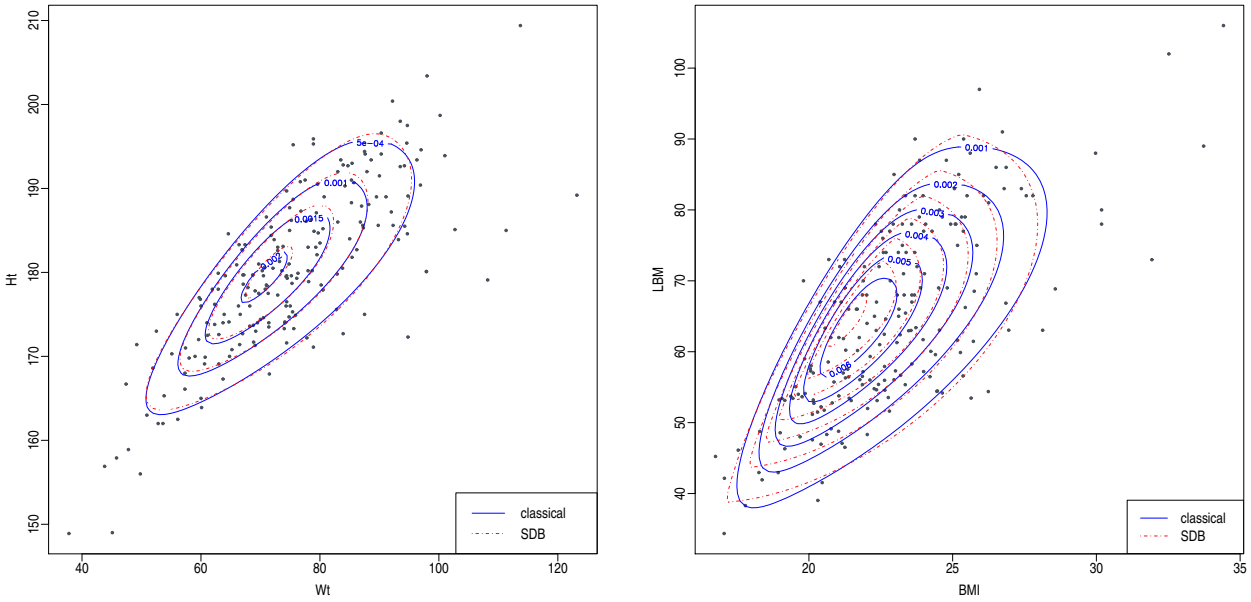


Figure 1: AIS data: scatter plot of (Wt, Ht) in the left panel and of (BMI, LBM) in the right panel, with superimposed fitted distributions using the classical and the SDB skew-normal families.

A qualitative comparison of the formal properties of the two distributions lends several annotations. Some of these have already been presented by Sahu et al. (2003), but they are included here for completeness.

1. Distributions (2) and (5) are regulated by the parameter sets $(\xi, \Omega, \alpha)$ and $(\xi, \Delta, \lambda)$. The two sets match as for nature and dimension of the three components; the number of individual parameter values is $2d + d(d+1)/2$ in both cases.

2. As remarked by Sahu et al. (2003), the two distributions coincide only for $d = 1$. We underline that for $d > 1$ neither one of the two families is a subset of the other. To see this fact, consider matching the arguments of $\exp(\cdot)$ in the moment generating functions by choosing a common $\xi$ and $\Omega = \Delta + \Lambda^2$. The final factors of $M_c$ and $M_s$, however, cannot match if $d > 1$, because in $M_s$ there is the product of several $\Phi$ terms, while $M_c$ has only one such term.

3. If $d$ is not small, computation of $f_s$ becomes cumbersome because of the $\Phi_d$ factor. Computational techniques, such as those described by Genz and Bretz (2009), facilitate handling this problem, but the burden involved increases rapidly with $d$. For example, LM2012 only consider examples with $d \leq 4$.

4. The classical skew-normal family is closed under affine transformations, while the same fact does not hold for the SDB family. The first part of this statement is widely documented in the above-quoted literature. For the second part, consider a linear transformation $X_s = AY_s$ for some non-singular matrix $A$, and notice that this is equivalent to replacing $\Lambda$ in (6) by a non-diagonal matrix.

5. Another remark of Sahu et al. (2003) is that $f_s$ can allow for $d$ independent skew-normal components, when $\Delta$ is diagonal, while $f_c$ can factorize only as a product where at most one factor is skew-normal with non-vanishing slant parameter.

6. Stochastic representations (3) and (6) involve 1 and $d$ latent variables, respectively. The latter one seems to fit less easily in an applied setting, because it requires that for each observed component there is a matching latent component, while the classical construction can more easily be incorporated in the logical frame describing a real phenomenon subject to selective sampling based on one latent variable.

7. For the classical skew-normal, the expressions of higher order cumulants and Mardia's coefficients of multivariate skewness and kurtosis are given in Appendix A.2 of Azzalini and Capitanio (1999). The range of skewness is $[0, g_1^*)$ where $g_1^* = 2(4-\pi)^2/(\pi-2)^3 \approx 0.9905658$; the range of excess kurtosis is $[0, g_2^*)$, where $g_2^* = 8(\pi-3)/(\pi-2)^2 \approx 0.869177$.

   For the SDB form, expressions of Mardia's coefficients are given in an appendix of the present paper. Numerical maximization of these expressions when $d = 2$ leads to the range $[0, 1.98113)$ and $[0, 1.738354)$, whose maximal values appear to coincide numerically with $2\, g_1^*$ and $2\, g_2^*$, respectively.

8. For the classical skew-normal, the distribution of quadratic forms can be obtained from the similar case under normality; for instance, we can state that $(Y_c - \xi)^\top \Omega^{-1} (Y_c - \xi) \sim \chi^2_d$. This fact is useful, among other things, for building QQ-plots of model adequacy. No similar result is known to hold for the SDB form.

Because the indication emerging from these remarks is not completely in favour of either of the two formulations, the choice is a matter of emphasis on one or another aspect, depending on the specific problem of interest or even individual preference.

However, for the Mardia's coefficients, we notice that the wider range of the SDB formulation is not a crucial aspect, because for both $f_c$ and $f_s$ these coefficients have a very small range anyway, and the limitations are to be overcome by the introduction of the corresponding skew-$t$ families. A possibly more important point is the factorization issue of Remark 5, if this aspect is relevant for the specific problem under consideration.

On the other hand, the lack of closure under affine transformations can be a non-negligible limitation in a range of problems. Sahu et al. (2003, p. 132) observe that the $f_s$ distribution could be extended to the case of non-diagonal matrices. This route does not seem to have been developed, presumably because it has been subsumed in the even more general 'closed skew-normal' formulation of Gupta et al. (2004) and González-Farías et al. (2004). The adoption of the closed skew-normal distribution would also have the advantage of replacing the factor $\Phi_d$ in (5) by one of type $\Phi_m$, where we can take $m < d$. Also the interpretability issue of Remark 6 can be of relevance if one is concerned with qualitative and subject-matter plausability of the formulation.

An additional aspect of these densities is illustrated graphically. Figure 2 refers to bivariate distributions having parameters

$$\xi = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad \Omega = \Delta = \begin{pmatrix} 1 & 1/2 \\ 1/2 & 1 \end{pmatrix}, \qquad \alpha = \lambda = \begin{pmatrix} 30 \\ 50 \end{pmatrix}$$

with $(\Omega, \alpha)$ employed for the clasical SN on left-side plot and $(\Delta, \lambda)$ employed for the SDB density on the right-hand plot. The slant parameters $\alpha$ and $\lambda$, respectively, are sufficiently large to indicate the limit behaviour of the densities when these parameters diverge. In the first case, the support of the density approaches the half-plane $\{(x, y) : \alpha_1 \omega_1^{-1} x + \alpha_2 \omega_2^{-1} y \geq 0\}$, with $\omega_1 = \omega_2 = 1$ in this case; in the second diagram, the support reduces to one of the quadrants, identified by the signs of the $\lambda$ components.

A well-known issue with the likelihood function associated to the classical SN distribution is the systematic presence of a stationarity point of the profile log-likelihood at $\alpha = 0$. This aspect is linked to another one, that is, singularity of the information matrix at $\alpha = 0$. Given the ensuing complication to the inferential

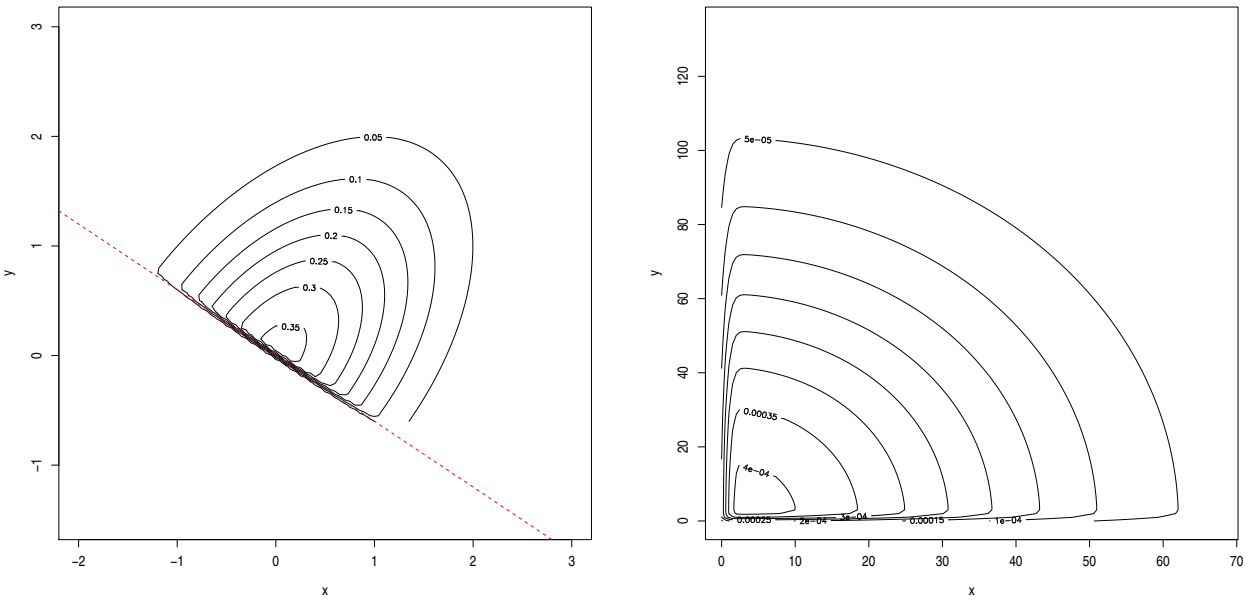


Figure 2: Contour level plots of the classical and SDB density where the slant parameters diverge.

process when the MLE estimate $\hat{\alpha}$ is close to 0, a natural question is whether the SDB version is free of this problem.

A simple answer is possible in the special case when $\Delta$ is diagonal. Factorization of the $f_s$ density, underlined in Remark 5, translates into factorization of the likelihood function into the product of $d$ likelihoods, each of univariate SN type with slant parameter equal to 0. Hence the stationarity problem and its implication persist also with the SDB density.

When $\Delta$ is not diagonal, the question becomes algebraically more intricate. Because a detailed analysis of the question would involve a project on its own, we have confined ourselves to some numerical explorations, which have confirmed the qualitative indications of the case with a diagonal $\Delta$. The end conclusion is that, also on this front, the two formulations behave in a qualitatively analogous way.

## 2.2 Corresponding skew-$t$ distributions

Each of the two skew-normal families discussed above leads to a matching form of skew-$t$ family. For the classical case, this can be obtained by replacing the assumption of joint normality of $(X_0, X_1)$ in (3) by one of $(d+1)$-dimensional Student's $t$ distribution; this has been formulated by Branco and Dey (2001) and examined further by Azzalini and Capitanio (2003). The SDB skew-$t$ has been obtained by Sahu et al. (2003) assuming that $(Z, \varepsilon)$ entering (6) is a $(2d)$-dimensional Student's $t$.

In both cases, the resulting density is similar in structure to the skew-normal case, with the $\varphi_d$ term replaced by a $d$-dimensional $t$ density on $\nu$ degrees of freedom, but the skewing factor is different: for the classical version, it is given by the distribution function of a univariate $t$ on $\nu + d$ degrees of freedom; for the SDB version, the $t$ distribution function is $d$-dimensional. The shape of these densities is qualitatively similar to their corresponding SN density, but with more elongated tails; therefore, we do not present specific plots.

The classical skew-$t$ can also be generated via a construction that closely resembles that of the regular Student's $t$, as follows: given a classical skew-normal variate $Z_c$ with location $\xi = 0$ and an independent

$V \sim \chi^2_\nu/\nu$, the variable

$$X_c = \xi + Z_c/\sqrt{V}$$

has the same distribution as the classical skew-$t$ described above. We are not aware of a similar type of representation for the SDB skew-$t$.

Both types of skew-$t$ distribution involve only one extra parameter, $\nu$, compared to the skew-normal. However, this is sufficient to widen considerably the range of feasible values of the coefficients of skewness and kurtosis, both marginally for each component and for the global Mardia's coefficients. It also allows to regulate separately skewness and kurtosis, which are instead interlinked in the skew-normal distribution. We refer the reader to the original sources for more details on this point. This increased flexibility in the skew-$t$ family, in either version, makes it especially appealing for practical data fitting; we shall explore this aspect in the next section.

As anticipated earlier, we now discuss some of the statements by Lee and McLachlan, in the paper denoted earlier LM2012 and related ones, with respect to the formal properties established above; an analogous discussion concerning the more applied side is deferred to the subsequent sections. In these expositions the skew-normal and skew-$t$ distributions which we refer to as classical are named 'restricted' and those that we refer to as SDB are named 'unrestricted', under the incorrect presumption that the SDB variants constitute a more general family than the classical ones.

To be specific, the opening sentence of Section 2.2 of LM2012 states that "The unrestricted multivariate skew-normal (uMSN) distribution can be viewed as a simple extension of the rMSN distribution in which the univariate latent variable $U_0$ is replaced by a multivariate analogue, that is, $U_0$", where rMSN stands for restricted multivariate skew-normal and their $U_0$ is $V_1$ in (4). The same conviction is replicated near the end of the same section.

Not only does the adopted terminology, restricted vs. unrestricted, convey a message of broader generality in the second case, but this is also explicit in the use of the term 'extension', which is inappropriate because we have seen earlier that neither one of the two families is a subset of the other for $d > 1$. Notice that no such claim of wider generality had been made by Sahu et al. (2003).

Furthermore, in the discussion of the classical skew-$t$ distribution near the end of Section 3 of LM2012, the authors state that "the form of skewness is limited in these characterizations. In Sect. 5 we study an extension of their approach to the more general form of skew $t$-density as proposed by Sahu et al. (2003)." This claim of limited form of skewness is completely unsupported; no expressions of any measure of skewness is even reported. As a matter of fact, the opposite is true: the range of skewness for the classical skew-$t$ distribution is unlimited both marginally, when measured by the usual coefficient $\gamma_1$, and globally, when measured by Mardia's coefficient $\gamma_{1,d}$. To see this fact in the univariate case, which coincides with the behaviour of univariate components in a multivariate skew-$t$ distribution, consider the expression of $\gamma_1$ on p. 382 of Azzalini and Capitanio (2003) and let $\nu \to 3$; for the Mardia's coefficient, see (6.31) on p. 178 of Azzalini and Capitanio (2014).

# 3 Model-based Clustering with Skew-$t$ Distributions

## 3.1 Preliminary remarks

The above-recalled distributions have been employed not only in a range of numerical illustrations of theoretical papers but also in various genuine applied problems and as ingredients for methodological developments. For the purpose of data-fitting, the skew-$t$ distributions, in either form, have more potential because of the additional tail-weight parameter. Although in both variant forms there is a single parameter regulating the tail behaviour, this has proved to work effectively in a large number of cases.

It is not our purpose here to present a general review of all this work, which would take considerable space, but at least we like to mention the paper of Azzalini and Genton (2008) because it provides numerical illustrations in several diverse areas, with favourable results in all cases. Taking into account that the SDB

formulation seems to have been employed relatively more often in the context of mixture distributions of type (1), we focus our attention in this direction by focusing on model-based clustering applications.

Moreover, model-based clustering represents the context where LM2012 and related publications are placed. Again, some of the statements made therein have provided further motivation for writing the present note. Similarly to the questionable statements on the theoretical aspects discussed in Section 2, LM2012 contains claims like "the superiority of FM-uMST model is evident" (end of Section 6.1) and "the extra-flexibility offered by the more general FM-uMST model" (end of Section 6.2). These comments refer to two specific cases, but are taken as the basis for the general statement "Examples on several real datasets shows [sic] that the unrestricted model is capable of achieving better clustering than the restricted model" (Section 8).

Besides these specific passages, the general tone of the paper conveys a message of overall superiority of the SDB formulation, and we believe this is unjustified. It is the purpose of the remaining part of this section to explore in a more systematic and fairer way the performance of the two forms of skew-$t$ distributions in clustering applications. An important aspect in these comparisons is the potential bias due to selective reporting. To avoid the perception of bias that can arise from selection of a subset of variables from a data set like the AIS data, we examine all possible pairs and triplets of variables in our real data illustrations (Section 3.5).

Note that there is much other work on model-based clustering using skewed distributions. Some of this work is based on the formulations already discussed in this paper (e.g., Pyne et al., 2009; Lin, 2009, 2010; Lin et al., 2013; Vrbik and McNicholas, 2012, 2014). However, there is also a substantial amount of work on model-based clustering using other skewed distributions (e.g., Karlis and Santourian, 2009; Ho et al., 2012; Murray et al., 2013; McNicholas et al., 2013; Browne and McNicholas, 2013; Franczak et al., 2014). In the analyses given herein, we focus only on comparison of the two skew-$t$ formulations that are the subject of the present paper.

### 3.2 Performance assessment

Although the illustrations herein are conducted as genuine cluster analyses, i.e., no knowledge of labels is assumed, the true labels are known; accordingly, we can asses the classification performance of the fitted models. The adjusted Rand index (ARI; Hubert and Arabie, 1985), which is a corrected version of the Rand index (Rand, 1971), is used for this purpose. The ARI takes a value of one when there is perfect agreement between two classes, or two partitions in general, and its expected value is zero under random classification. Negative values are also possible and indicate classification that is worse than would be expected by guessing.

### 3.3 Implementation

We implement the classical and SDB skew-$t$ mixture models in the R software (R Core Team, 2013), and our implementations will be made available as a function within the `mixture` package. Starting values were obtained using $k$-means clustering and for each data set, the same starting values were used for each approach.

### 3.4 Simulated data

A two-dimensional data set is simulated to illustrate a difficult clustering problem, where one cluster lies within the other (Figure 3). One component (black in Figure 3) is generated from a SDB skew-$t$ distribution and the other (red in Figure 3) from a classical skew-$t$ distribution. Each component contains 1,000 observations. Mixtures of classical and SDB skew-$t$ distributions, respectively, are fitted to these simulated data in a model-based clustering framework with $G = 2$ components. The results (Figure 4 and Table 1) show that both classical and SDB skew-$t$ mixtures give very good classification performance, with the classical

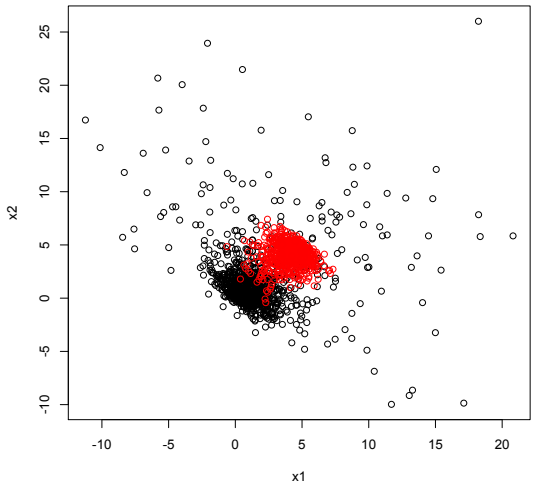


Figure 3: Scatter plot illustrating the simulated data, coloured by component.

skew-$t$ mixture giving slightly better classification performance (ARI=0.781) than the SDB skew-$t$ mixture (ARI=0.774).

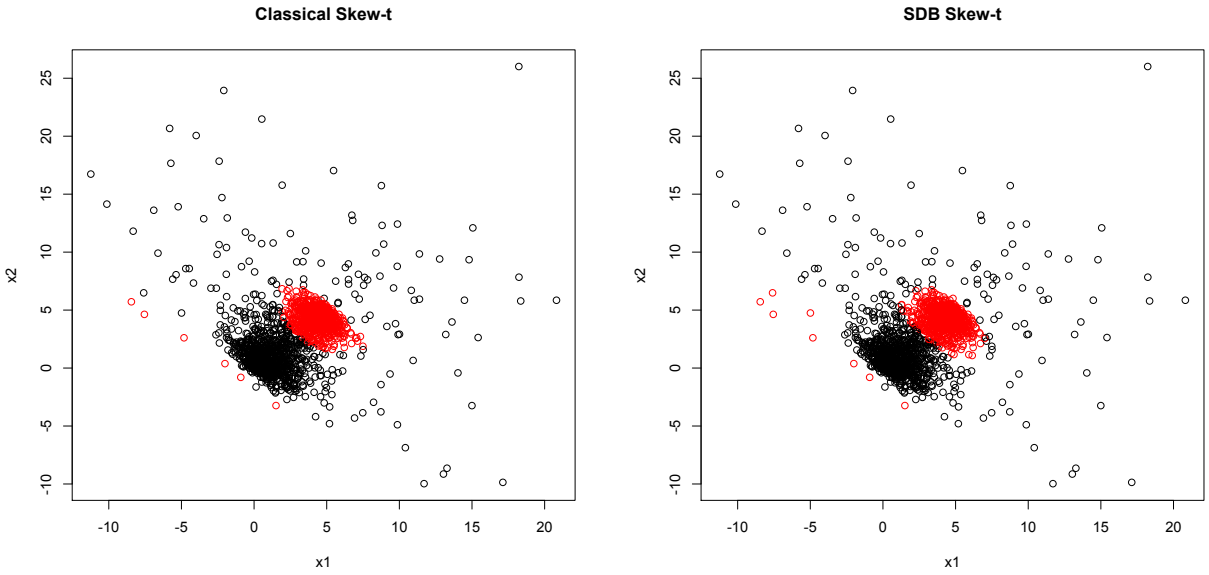


Figure 4: Scatter plot of the predicted classifications for the classical (left) and SDB (right) skew-$t$ mixtures applied to the simulated data, where colour represents predicted classifications.

## 3.5 Crabs data

Campbell and Mahon (1974) report data on five biological measurements of 200 crabs of genus *leptograpus*. The data were collected in Fremantle, Western Australia, and comprise 50 male and 50 female crabs for each

Table 1: Cross-tabulations of true (A,B) and predicted (1,2) group memberships for the classical and SDB skew-$t$ mixtures, respectively, for the simulated data.

| | Classical | | SDB | |
|---|---|---|---|---|
| | 1 | 2 | 1 | 2 |
| A | 952 | 48 | 942 | 58 |
| B | 68 | 932 | 62 | 938 |

of two species. These data are available in the `MASS` package for R, which contains data sets from Venables and Ripley (1999). We focus on these data because they have been used throughout the model-based clustering literature to illustrate the performance of various methods.

First, we focus on two of the five variables, i.e., frontal lobe size (FL) and rear width (RW), both measured in millimetres. Based on these variables, it should be possible to recover gender within a clustering framework (cf. Figure 5). Mixtures of classical and SDB skew-$t$ distributions, respectively, are fitted to the crabs data in a model-based clustering framework with $G = 2$ components. The results (Table 2 and Figure 6) show

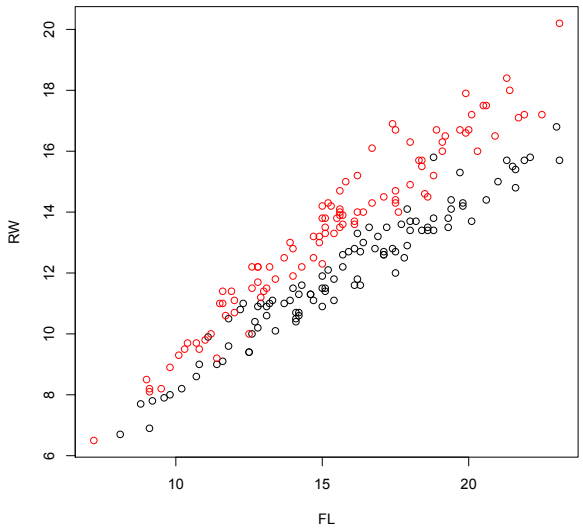


Figure 5: Scatter plot illustrating two variables (FL, RW) from the crabs data, coloured by gender.

that the classical skew-$t$ mixture model gives good classification performance with only 30 misclassifications, corresponding to an ARI of 0.487. The SDB skew-$t$ mixture model, however, gives very poor classification, producing results only slightly better than would be expected under random classification (ARI=0.074), with 72 of the 200 crabs misclassified.

Table 2: Cross-tabulations of true and predicted group memberships for the classical and SDB skew-$t$ mixtures, respectively, for the crabs data.

| | Classical | | SDB | |
|---|---|---|---|---|
| | 1 | 2 | 1 | 2 |
| Male | 83 | 17 | 83 | 17 |
| Female | 13 | 87 | 55 | 45 |

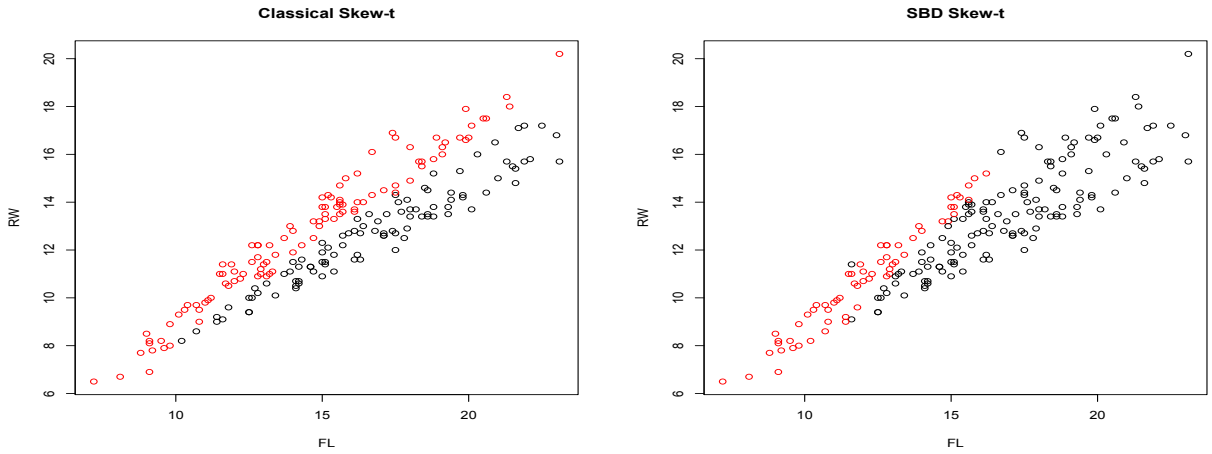


Figure 6: Scatter plot illustrating predicted classifications for the classical (left) and SDB (right) skew-$t$ models applied to the crabs data, where colour represents predicted classifications.

We previously remarked (Section 3.1) that presenting results based on a subset of variables in a data set may incur bias or at least a perception thereof. Accordingly, a more comprehensive analysis is also undertaken, wherein mixtures of classical and SDB skew-$t$ distributions, respectively, are applied to all possible pairs and triplets of variables of the crabs data within a model-based clustering framework. Note that we take gender as the true class, which is consistent with the analysis of Peel and McLachlan (2000) who considered two variables of the crabs data. The results of these analysis are presented in Figure 7, which gives the ARI values for the analyses with the classical and SDB skew-$t$ distributions, respectively, for all pairs and triplets. It is visible that in most cases the points are close to the identity line; hence indicating no advantage of either one formulation over the other. The only point departing significantly from this pattern is based on a pair and indicates a preference for the classical skew-$t$ distribution; this point corresponds to the pair (FL, RW) examined earlier. The group of points near the bottom-left corner arises because some of the variables do not exhibit any discriminatory capability due largely to their very high correlation.

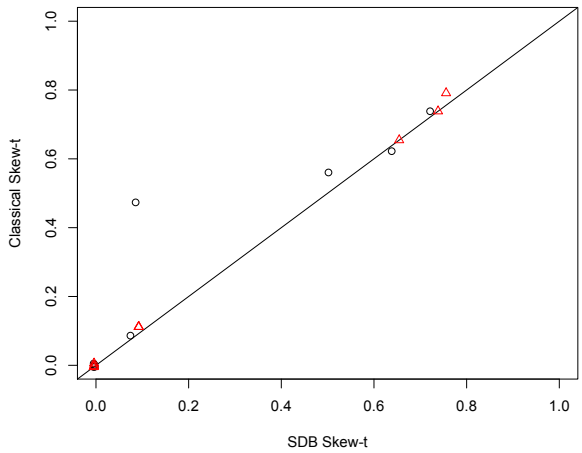


Figure 7: ARI values from analyses of all pairs (black circles) and triplets (triangles) of variables in the crabs data using classical and SDB skew-$t$ distributions, respectively.

### 3.6 AIS data

We also consider the AIS data, which provide 11 biomedical and anthropometric measurements and two categorial variables, i.e., gender and sport, for each of 202 Australian athletes. Some of these variables have already appeared herein (Figure 1). The reason for considering these data is that they have been extensively used for illustrative purposes in the literature on skew-symmetric distributions.

We proceed similarly to the crabs data, by considering all possible pairs and triplets of the continuous measurements to build clusters of the data points. To compute the ARI index, we again take gender as the reference classifying variable. The results are summarized in Figure 8, which is constructed in an analogous fashion to Figure 7. The message clearly delivered by the new figure is that of an overall equivalence between the two families of distributions in this comparative exercise.

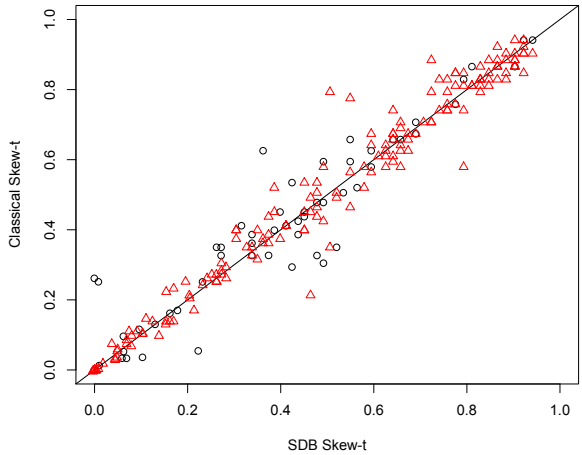


Figure 8: ARI values from analyses of all pairs (black circles) and triplets (triangles) of variables in the AIS data using classical and SDB skew-$t$ distributions, respectively.

### 3.7 Comments

One may argue that our results for the analysis of the crabs and AIS data sets are merely an artifact of our software. Therefore, we repeated the analyses using the R packages `EMMIXskew` (Wang et al., 2013) and `EMMIXuskew` (Lee and McLachlan, 2013a), which implement the 'restricted' and 'unrestricted' skew-$t$ distributions, respectively, borrowing the authors' terminology. The results for the crabs data (Figure 9) are are similar to ours (Figures 7 and 8), but with more dispersion around the identity line. The results for the AIS data also exhibit more variability around the identity line but we also note that there are clearly many more points above the identity line than below it, i.e., there are many more pairs and triples for which the classical formulation outperforms the SDB formulation (cf. Figure 9).

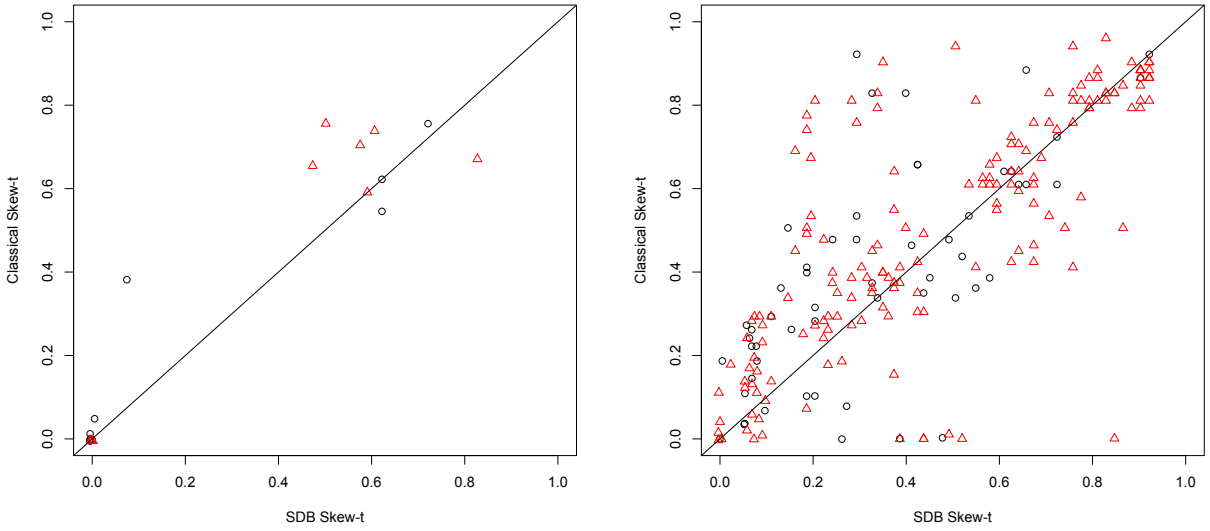


Figure 9: ARI values from analyses in Sections 3.5 (left) and 3.6 (right) repeated using the `EMMIXskew` and `EMMIXuskew` packages; for both the crabs (left) and AIS (right) results, pairs are represented by circles and triplets by triangles.

The extra variability around the identity line in the results using these R packages (Figure 9) over ours (Figures 7 and 8) can be attributed to the better control on the starting values with our R code. Specifically, we used the same starting values for the classical and SDB skew-$t$ distributions in each case; however, this was not possible through the `EMMIXskew` and `EMMIXuskew` packages.

## 4 Conclusion

We have presented various arguments, of mathematical and of numerical nature, to refute a claimed superiority of the SDB construction over the classical one within the domain of skew-symmetric distributions.

An argument made throughout this paper is that the nomenclature "restricted" and "unrestricted" is inappropriate. The classical skew-$t$ is not a restricted version of the other neither from the formal viewpoint nor in the observed numerical outcomes, both with simulated and with real data. As for their effectiveness in numerical work, there is little to choose between the two formulations and, if one really wants to rank them, the only case showing some difference was with the crabs data, where the classical version performed slightly better.

On the more general side, the points examined in the qualitative comparison of Section 2.1 indicate that neither formulation is uniformly superior and the choice between them depends on the emphasis on one

or another aspect. In our opinion, the property of closure under affine transformations (Remark 4) and, from the interpretability viewpoint, the more plausible stochastic representation (Remark 6) are of greater general importance than the possibility of factorization in $d$ independent components (Remark 5), leading to some preference for the classical skew-$t$, although this is not a universal statement and in specific cases this inclination might be reversed.

## Acknowledgements

We are grateful to Márcia Branco for fruitful discussion of various aspects of the SDB formulation and for making available to us related material.

## References

Azzalini, A. and A. Capitanio (1999). Statistical applications of the multivariate skew normal distribution. *J. Roy. Stat. Soc,, series B 61*(3), 579–602. Full version of the paper at `arXiv.org:0911.2093`.

Azzalini, A. and A. Capitanio (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew $t$ distribution. *J. Roy. Stat. Soc,, series B 65*(2), 367–389. Full version of the paper at `arXiv.org:0911.2342`.

Azzalini, A. with the collaboration of A. Capitanio (2014). *The Skew-Normal and Related Families*. IMS monographs. Cambridge University Press.

Azzalini, A. and A. Dalla Valle (1996). The multivariate skew-normal distribution. *Biometrika 83*, 715–726.

Azzalini, A. and M. G. Genton (2008). Robust likelihood methods based on the skew-$t$ and related distributions. *Int. Stat. Rev. 76*, 106–129.

Branco, M. D. and D. K. Dey (2001). A general class of multivariate skew-elliptical distributions. *J. Multivariate Analysis 79*, 99–113.

Browne, R. P. and P. D. McNicholas (2013). A mixture of generalized hyperbolic distributions. arXiv:1305.1036.

Campbell, N. A. and R. J. Mahon (1974). A multivariate study of variation in two species of rock crab of genus leptograpsus. *Australian Journal of Zoology 22*, 417–425.

Franczak, B. C., P. D. McNicholas, R. P. Browne, and P. M. Murray (2014). Parsimonious shifted asymmetric Laplace mixtures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. To appear.

Genz, A. and F. Bretz (2009). *Computation of multivariate normal and t probabilities*. Springer-Verlag.

González-Farías, G., J. A. Domínguez-Molina, and A. K. Gupta (2004). Additive properties of skew normal random vectors. *J. Stat. Plann. & Inference 126*, 521–534.

Gupta, A. K., G. González-Farías, and J. A. Domínguez-Molina (2004, April). A multivariate skew normal distribution. *J. Multiv. Analysis 89*(1), 181–190.

Ho, H. J., S. Pyne, and T. I. Lin (2012). Maximum likelihood inference for mixtures of skew Student-t-normal distributions through practical EM-type algorithms. *Statistics and Computing 22*(1), 287–299.

Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of Classification 2*, 193–218.

Karlis, D. and A. Santourian (2009). Model-based clustering with non-elliptically contoured distributions. *Statistics and Computing 19*(1), 73–83.

Lee, S. and G. J. McLachlan (2014). Finite mixtures of multivariate skew $t$-distributions: some recent and new results. *Statistics and Computing 24*(2), 181–202. Published online 20 October 2012.

Lee, S. X. and G. J. McLachlan (2013a). EMMIXuskew: An R package for fitting mixtures of multivariate skew $t$ distributions via the EM algorithm. *Journal of Statistical Software 55*(12), 1–22.

Lee, S. X. and G. J. McLachlan (2013b). Model-based clustering and classification with non-normal mixture distributions. *Stat. Methods & Applications 22*, 427–454.

Lee, S. X. and G. J. McLachlan (2013c). On mixtures of skew normal and skew $t$-distributions. *Advances in Data Analysis and Classification 7*, 241–266.

Lin, T.-I. (2009). Maximum likelihood estimation for multivariate skew normal mixture models. *Journal of Multivariate Analysis 100*, 257–265.

Lin, T.-I. (2010). Robust mixture modeling using multivariate skew $t$ distributions. *Statistics and Computing 20*(3), 343–356.

Mardia, K. (1970). Measures of multivariate skewness and kurtosis with applications. *Biometrika 57*, 519–530.

Lin, T.-I., G. J. McLachlan, and S. X. Lee (2013). Extending mixtures of factor models using the restricted multivariate skew-normal distribution. arXiv:1307.1748.

McNicholas, S. M., P. D. McNicholas, and R. P. Browne (2013). Mixtures of variance-gamma distributions. Arxiv preprint arXiv:1309.2695.

Murray, P. M., R. P. Browne, and P. D. McNicholas (2013). Mixtures of skew-$t$ factor analyzers. arXiv:1305.4301v2.

Peel, D. and G. J. McLachlan (2000). Robust mixture modelling using the t distribution. *Statistics and Computing 10*(4), 339–348.

Pyne, S., X. Hu, K. Wang, E. Rossin, T.-I. Lin, L. M. Maier, C. Baecher-Allan, G. J. McLachlan, P. Tamayo, D. A. Hafler, P. L. De Jager, and J. P. Mesirow (2009). Automated high-dimensional flow cytometric data analysis. *Proceedings of the National Academy of Sciences 106*, 8519–8524.

R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association 66*(336), 846–850.

Sahu, K., D. K. Dey, and M. D. Branco (2003). A new class of multivariate skew distributions with applications to Bayesian regression models. *Canad. J. Statist. 31*(2), 129–150. Corrigendum: vol. 37 (2009) , 301–302.

Venables, W. N. and B. D. Ripley (1999). *Modern Applied Statistics with S-PLUS*. Springer.

Vrbik, I. and P. D. McNicholas (2012). Analytic calculations for the EM algorithm for multivariate skew-mixture models. *Statistics and Probability Letters 82*(6), 1169–1174.

Vrbik, I. and P. D. McNicholas (2014). Parsimonious skew mixture models for model-based clustering and classification. *Computational Statistics and Data Analysis 71*, 196–210.

Wang, K., A. Ng, and G. McLachlan. (2013). *EMMIXskew: The EM Algorithm and Skew Mixture Distribution*. R package version 1.0.1.

# Appendix: Cumulants and Mardia's coefficients for SDB skew-normal

For the SDB skew-normal, differentiation of $K_s(t) = \log M_s(t)$ produces

$$\nabla K_s(t) = \xi + (\Delta + \Lambda^2)t + [\zeta_1(\lambda_j t_j)\lambda_j]_{j=1}^d$$
$$\nabla\nabla^\top K_s(t) = (\Delta + \Lambda^2) + \mathrm{diag}(\zeta_2(\lambda_1 t_1)\lambda_1^2, \ldots, \zeta_2(\lambda_d t_d)\lambda_d^2),$$

where $\zeta_r(x)$ is the $r$th derivatives of $\zeta_0(x) = \log\{2\,\Phi(x)\}$. Evaluation at $t = 0$ gives (8).

Further differentiation and evalutation at 0 gives the 3rd order cumulant

$$\kappa_{rst} = \begin{cases} \zeta_3(0)\,\lambda_r^3 & \text{if } r = s = t, \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

where

$$\zeta_3(0) = b\,(4/\pi - 1) = (2/\pi)^{3/2}\,(4 - \pi)/2\,.$$

We can now compute the Mardia's coefficient $\gamma_{1,d}$ of multivariate skewness, recalling that the 3rd order cumulant coincides with the 3rd order central moment. Denote by $\Sigma = (\sigma_{rs})$ the variance matrix in (8) and let $\Sigma^{-1} = (\sigma^{rs})$, $\mu_j = b\lambda_j$. From (2.19) of Mardia (1970), write

$$\begin{aligned} \gamma_{1,d} &= \sum_{rst}\sum_{r's't'} \kappa_{rst}\kappa_{r's't'}\sigma^{rr'}\sigma^{ss'}\sigma^{tt'} \\ &= \zeta_3(0)^2 \sum_{u,v} \lambda_u^3\,\lambda_v^3\,(\sigma^{uv})^3 \\ &= \left(\frac{4-\pi}{2}\right)^2 \sum_{u,v} \mu_u^3\,\mu_v^3\,(\sigma^{uv})^3 \\ &= \left(\frac{4-\pi}{2}\right)^2 (\mu^{(3)})^\top \Sigma^{(-3)} \mu^{(3)}, \end{aligned} \tag{10}$$

where $\mu^{(3)}$ is the vector with elements $\mu_j^3$ and $\Sigma^{(-3)} = \big((\sigma^{uv})^3\big)$.

For (10) we do not have an expression of the maximal value. Numerical exploration indicates that the maximal value is 1.98113, that is, about the double value of the classical SN. This maximal value of $\gamma_{1,d}$ is obtained, irrespectively of $\Delta$, in these four cases:

$$\lambda = h\begin{pmatrix} \pm 1 \\ \pm 1 \end{pmatrix}, \quad \text{when } h \to \infty\,. \tag{11}$$

Derivation of the 4th order cumulants is similar to (9), leading to

$$\kappa_{rstu} = \begin{cases} \zeta_4(0)\,\lambda_r^4 & \text{if } r = s = t = u, \\ 0 & \text{otherwise,} \end{cases} \tag{12}$$

where

$$\zeta_4(0) = 2\,(\pi - 3)\,(2/\pi)^2 \approx 0.114771.$$

From here the Mardia's coefficient of (excess) kurtosis is

$$\begin{aligned}
\gamma_{2,d} &= \sum_{rstu} \kappa_{rstu}\sigma^{rs}\sigma^{tu} \\
&= \zeta_4(0)\sum_u \lambda_u^4\,(\sigma^{uu})^2 \\
&= 2\,(\pi-3)\sum_u \mu_u^4\,(\sigma^{uu})^2 \\
&= 2\,(\pi-3)(\mu^{(2)})^\top (I_d \odot \Sigma^{-1})^2 \mu^{(2)}, \qquad (13)
\end{aligned}$$

where $\mu^{(2)} = (\mu_1^2, \ldots, \mu_d^2)^\top$ and $\odot$ is the Hadamard product. A numerical search indicates that the maximal value of $\gamma_{2,d}$ is again achieved, irrespectively of $\Delta$, with $\lambda$ as in (11). The maximal observed value of the coefficient is 1.7383546, again twice the corresponding value in the classical skew-normal case.